CHAPTER 15

Distribution

We have emphasized that ecological economics is concerned with three issues: the allocation of resources, their distribution, and the scale of the economy. We have seen how the ecological sustainability of the Earth is related to the size or scale of the macroeconomy. We have also explored the economist's meaning of efficient allocation in our discussion of microeconomics and the basic market equation. We then looked at the macroeconomic allocation problem in Chapter 16. But the second issue, distribution and the fairness thereof, has remained largely in the background.

Pareto Optimality

In dealing with allocation, we saw that economics defines efficiency as the Pareto optimal allocation of resources by the market. This definition assumes a given distribution of wealth and income. More specifically, an efficient allocation is one that best satisfies individual wants *weighted by the individual's ability to pay*—that is, by her income and wealth. Change the distribution of income and wealth, and we get a different set of efficient prices (since different people want different things), which define a different Pareto optimum. Because different Pareto optima are based on different distributions of income and wealth, economists are reluctant to compare them; one optimum is as good as another. We saw that a major reason for scale expansion—economic growth—has been to avoid the issue of distributive equality. As long as everyone is getting more from aggregate growth, then the distributive issue is less pressing, at least as a cure for poverty. Besides, the efficiency of the allocation of aggregate growth loses its well-defined meaning (Pareto optimality) once we accept the legitimacy of changing distribution in the interest of fairness. Consequently, economics has tended to address distribution out of logical necessity, but quickly sets it aside in the interests of political convenience.

Box 15-1 Does a Pareto Optimal Allocation Assume a Given Scale as well as a Given Distribution?

If we take the concept of scale literally, as in the scale model of a house, to involve only a proportional change in all linear (scalar) dimensions, then we might say that a scale change is simply an increase or decrease in which all proportions remain constant. All relative prices, measuring unchanged relative scarcities, would also remain constant, defining an unchanging Pareto optimal allocation in terms of proportions. This seems to be what standard economists often have in mind. But is it possible to have everything grow in proportion? No, for two reasons. First, if something is fixed, it obviously cannot grow proportionally to everything else. What is fixed from the ecological economist's perspective is the size of the total ecosystem. As the economic subsystem grows, albeit proportionally in terms of its internal dimensions, the ecosystem itself does not grow. The economy becomes larger as a proportion of the total system—what we have called an increase in its scale, meaning size relative to the ecosystem. Natural capital becomes more scarce relative to manmade capital.

Of course, if the economy were to expand to encompass Earth's entire ecosystem (the model of "economic imperialism" in Chapter 3), the scale issue would disappear. In this sense the neoclassical economist's claim that if only all externalities were perfectly internalized then prices would automatically solve the scale problem (in the process of allocating everything in creation) makes sense. But it's a rather utopian point—like Archimedes' boast that he could move the Earth, if only he had a fulcrum and a long enough lever!

The second difficulty, long noticed by biologists and some economists, is that if you scale up anything (increase all linear dimensions by a fixed factor), you will inevitably change the relative magnitudes of nonlinear dimensions. Doubling length, width, and height will not double area; it will increase area by a factor of four and volume by a factor of eight. Biologists have long noted "the importance of being the right size." If a grasshopper were scaled up to the size of an elephant, it could not jump over a house. It would not even be able to move, because its weight (proportional to volume) would have increased eightfold, while its strength (proportional to a cross-sectional area of muscle and bone) would have increased only fourfold.

Returning to our example of a house, doubling the scale will increase surfaces and materials by fourfold and volumes to be heated or cooled by eightfold. Relative scarcities and relative prices cannot remain the same.

The answer to our question, does the notion of Pareto optimal allocation assume a given scale as well as a given distribution, appears to be yes. Size cannot increase proportionally because (1) there is a fixed factor, namely the size of the ecosystem, and (2) it is mathematically impossible even for all relevant internal dimensions of the subsystem to increase in the same proportion. In sum, it seems quite true that an optimal allocation assumes a given scale, just as it assumes a given distribution.

Economics prides itself on being a "positive science." Allocative efficiency is thought to be a positive, or empirically measurable, issue, even though, as we just saw, it presupposes a given distribution. Whether or not the scale of the economy is sustainable is also considered to be a positive issue involving biophysical constraints, although normative questions of conservation for the future and other species are not far below the surface. Distributive equity, on the other hand, is a normative issue. This is the main question addressed to distribution: Is it fair, not is it efficient, or is it ecologically sustainable? The question "Is it fair?" is directly and unavoidably normative, and for that reason alone given minimal attention by the positivist tradition of economics.

But like other sciences, economics assumes certain cultural values. First, the very criterion of objective efficiency, Pareto optimality, embodies an implicit normative judgment—namely, that malevolence or invidious satisfactions are not acceptable. If everyone but you becomes better off and you remain no worse off, the Pareto criterion tells us that is an objective increase in social welfare. But if everyone else is better off except you, and you are an envious person, then you will be less happy than before, even though your absolute situation is no worse. Economists must either make the (false) positive judgment that people are in fact not invidious and jealous, or the (true) normative judgment that envy at another's good fortune is a moral failing rather than a welfare loss.

There is a second reason that economics is less positive than some think. Redistribution can be efficient in the sense of increasing total social utility, yet economists make the value judgment that this kind of efficiency should not count. For example, redistributing a dollar from the low marginal utility use of the rich to the high marginal utility use of the poor will increase total utility to society and is in that sense efficient. The Pareto criterion forbids such interpersonal comparisons and summations of utility. Some argue that the major function of the Pareto criterion was precisely to sterilize the egalitarian implications of the law of diminishing marginal utility[1]—a law that economics cannot afford to give up, as we saw in our discussion of demand curves (see Chapter 9).

[1]J. Robinson, *Economic Philosophy*, Middlesex, England: Penguin, 1962.

If we admit interpersonal comparisons of utility, then distribution has efficiency implications as well as fairness implications. The extreme individualism of economics insists that people are so qualitatively different in their hermetical isolation one from another that it makes no sense to say that a leg amputation hurts Smith more than a pin prick hurts Jones. If we are all isolated individuals, we can rule out such obviously realistic human characteristics as envy and benevolence. Man as atomistic individual is the *Homo economicus* of neoclassical economics. Ecological economics' concept of the nature of man is "person-in-community," not isolated atom. Community here means community both with other humans and with the rest of the biosphere.

The Distribution of Income and Wealth

Ecological economics distinguishes between the distribution of income and the distribution of wealth, and between the functional and the personal distribution of income.

Wealth is a stock of assets, measured at a point in time, that is, cash in the bank, plus the market value of bonds, corporate shares, land, real estate, and consumer durables as of a given date. Income is the flow of earnings from these assets, plus the earnings of your own labor power (or human capital), between two dates, that is, over a period of time, usually a year. Labor power is not usually counted as capital because one cannot sell it all at once to another person (short of slavery) but can only rent it for certain durations. Income and wealth are thus two different magnitudes, measured in different units, and distributed differently over the population.[2] Wealth is usually more concentrated than income. And financial wealth is even more concentrated than wealth in general. In 1989, the top 1% owned 48% of financial wealth. Virtually all of the growth in wealth between 1983 and 1989 in the U.S. went to the top 20%. The bottom 80% was excluded from this growth, and the bottom 40% saw their wealth decline in absolute terms (Table 15.1).

Economics has a theory that explains income, as discussed next, and one that explains the prices of assets (though not entirely, as the "price" of entrepreneurship is a residual), but no theory at all to explain the distribution of wealth among individuals. It is the historical result of whose ancestors got there first, of marriage, of inheritance, plus individual ability and effort, and just plain luck.

[2]Wealth is measured in dollars (for example) and income in dollars/time. These magnitudes are as different as miles (distance) and miles per hour (speed).

Table 15.1

U.S. PERCENTAGE SHARE OF WEALTH AND INCOME BY PERCENTILE GROUP

	Percentile Shares			
Year	Top 1%	Next 19%	Bottom 80%	Gini Coefficient
Net Worth (Wealth)				
1983	33.8	47.6	18.7	0.799
1989	37.4	45.3	16.2	0.832
1992	37.2	46.6	16.3	0.823
1995	38.5	45.8	16.1	0.828
1998	38.1	45.3	16.6	0.822
Income				
1983	12.8	39.0	48.1	0.480
1989	16.4	39.0	44.5	0.521
1992	15.7	40.7	43.7	0.528
1995	14.4	40.8	44.9	0.518
1998	16.6	39.6	43.8	0.531

Source: E. N. Wolff, Top Heavy, The Twentieth Century Fund Report, *New York: New Press, 1995, p. 67 (years 1983–1992) and E. N. Wolff,* Recent Trends in Wealth Ownership, 1983–1998, *Working Paper No. 300, Table 2, Jerome Levy Economics Institute, April 2000.*

THE FUNCTIONAL AND PERSONAL DISTRIBUTION OF INCOME

Income distributed among people, regardless of its source, is called the personal distribution. Income is also distributed according to how much of total income goes to wages, interest, rent, and profit—the functional distribution. The idea behind the functional distribution is that income is not first created, then distributed. Rather, it is distributed as it is created among the factors combining to create it.

Remember from the circular flow diagram (see Figure 2.4) that supply and demand in the factors market determine the prices of factors—wages, interest, rent, with profit as a residual. Factor prices times the total amount of each factor used yields the functional distribution, usually expressed as percentage of total income going to landowners (rent), laborers (wages), capitalists (interest), and entrepreneurs (profit). Prices of each factor times the amount of the factor owned by each individual yields the personal distribution of income. The amount of each factor owned by each person, including labor power, is the personal distribution of wealth. Therefore, the personal distribution of wealth times the rental price of each type of wealth asset determines the personal distribution of income.

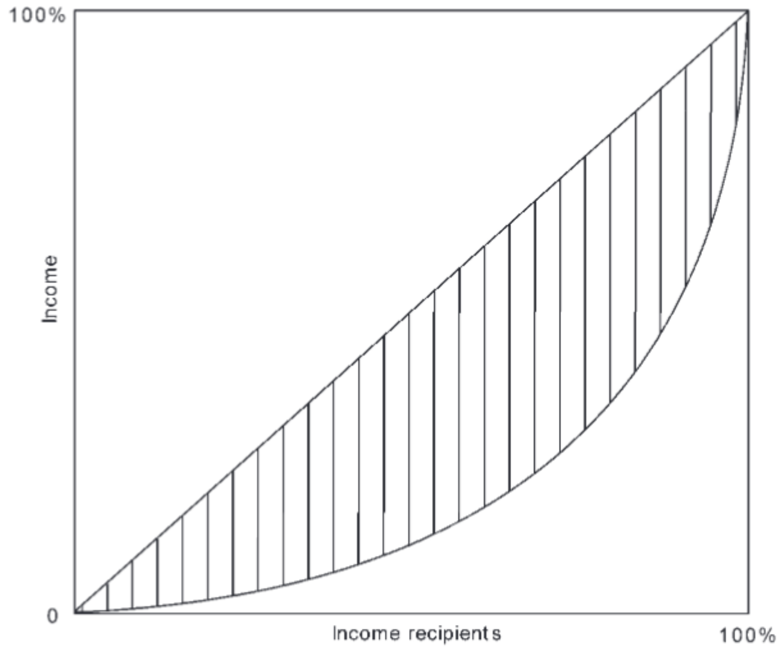

Figure 15.1 • The Lorenz curve. Because the Lorenz curve is in percentages, its shape does not depend on units of measure. It is therefore useful for making comparisons across countries and over time.

MEASURING DISTRIBUTION

Although economists have no good theory by which to explain the distribution of wealth and income, they do have useful ways of measuring and describing it statistically.[3] One useful representation is the **Lorenz curve**, shown in Figure 15.1. The X-axis shows the number of income recipients in terms of cumulative percentages, from lowest to highest income. The Y-axis shows the percentage of total income. The lengths of the axes are equal, so that when closed in, they make a square.

The Lorenz curve plots the percentage of total income going to each percentage of income recipient. We know that 0% of income recipients will get 0% of the income, and that 100% of income recipients will get 100% of the income, so we already know the two extreme points on any Lorenz curve. If each percentage of the population received the same percentage of the income (i.e., the bottom 20% got 20% of total income, the bottom 70% got 70% of income), we would have perfect equality. The Lorenz curve would be the 45-degree line connecting (0, 0) with (100, 100). But suppose the bottom 80% of recipients get 44% of the income.

[3] For a clear and insightful exposition, see J. Pen, *Income Distribution*, New York: Praeger, 1971.

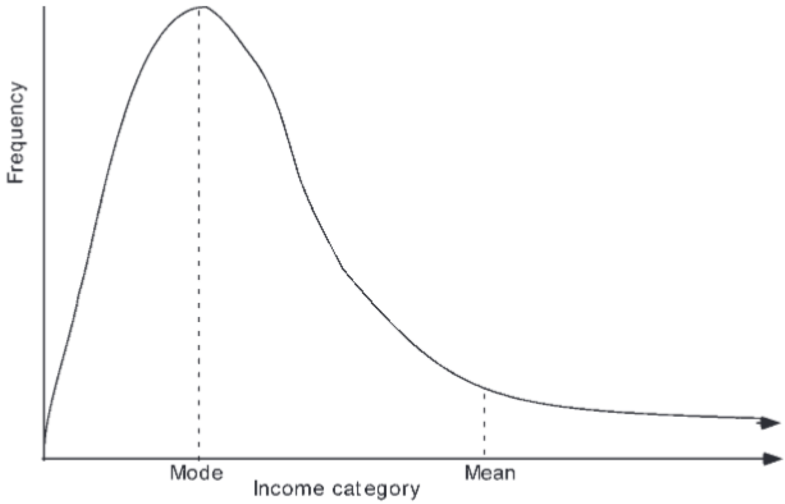

Figure 15.2 • The frequency distribution of income.

That gives us another point, one that lies well below the 45-degree line. If we fill in many points between the extremes, we get a curve shaped like the one in Figure 15.1. The closer the curve to the 45-degree line, the more equal the distribution; the farther away, the less equal. The shaded area defined by the curve and the 45-degree line measures inequality. In the limit, if one person got 100% of income and everyone else got 0%, the Lorenz curve would coincide with the axes and look like a backwards L.

The ratio of the shaded area (between the curve and the 45-degree line) to the total triangular area under the 45-degree line is called the **Gini coefficient**. For perfect equality the shaded area is zero, and consequently the Gini coefficient is 0; for perfect inequality the shaded area takes up the whole area under the 45-degree line, and consequently the Gini coefficient is 1. Values of the Gini coefficient for U.S. wealth and income distribution are given in Table 15.1.

The *Gini coefficient* is used to measure the inequality of the distribution of wealth or income across a population. A Gini coefficient of 1 implies perfect inequality (one person owns everything) and a coefficient of zero indicates a perfectly equal distribution.

A more familiar statistical description is the common frequency distribution, shown in Figure 15.2. The X-axis shows income category and the Y-axis shows number of members in each income category (frequency). Income distribution does not follow a normal distribution, as does height or many other personal characteristics. Rather, it is highly skewed, with the mode well below the mean, and a very, very long tail to the right needed to reach the top income.

If we wanted to show the maximum income on Figure 15.2, we would need a fold-out extending the horizontal axis by the length of a football field. Graphical representations generally do not capture the extreme inequality at the upper income range. Income categories are frequently truncated at a maximum category of "$100,000 and over," where "over" means four orders of magnitude over.

Moreover, these data are just for the United States. The distribution of wealth and income *between* countries is far greater than that found *within* countries.

What is the proper range of inequality in the distribution of income? Surely it is impossible to have one person owning everything, and everyone else owning nothing. Maybe we could have everyone else getting a subsistence wage, and the fortunate one person enjoying the entire social surplus above subsistence. But most people would not consider that fair, even though possible. At the other extreme, few people think a perfectly equal distribution—a Gini coefficient equal to zero—would be fair either. After all, some people work harder than others, and some jobs are more difficult than others. Fairness in a larger sense would require some income differences. There is a legitimate case to be made that differences in distribution provide a socially useful incentive for industriousness and innovation.

Is there a legitimate range of inequality, beyond which further inequality becomes either unfair or dysfunctional? What might such a range be? Plato thought that the richest citizen should be four times wealthier than the poorest. Ben Cohen and Jerry Greenfield, of Ben and Jerry's ice cream fame, at one time reportedly pledged that the highest paid executive would receive no more than five times the salary of the lowest paid employee. Maybe Plato, Ben, and Jerry were wrong, though, and maybe a factor of ten would be better. Or 20 or 50. Currently the acceptable ratio is not defined, and in 1999 in the United States, the *typical* CEO earned 475 times the *typical* worker.[4] (Of course, comparing the *top* CEOs to the *bottom* workers would yield even a higher ratio.) Ecological economics does not accept the current notion that real total output can grow forever. If the total is limited, then the maximum for one person is implicitly limited. The issue of a proper range of inequality in distribution is therefore critical for ecological economics, even though it has not yet received due attention. The standard economist's effort to keep distribution at bay forever by eternal growth is not a satisfactory solution.

Finally a word on the functional distribution of income. For industrial countries, the division varies around the following: wages = 70%, profits = 20%, interest = 8%, and land rent = 2%. For ecological economics, what is striking is that essentially none of the value of the total product is at-

[4] J. Reingold and F. Jesperson, Executive Pay: It Continues to Explode—and Options Alone Are Creating Paper Billionaires, *Business Week*, April 17, 2000. Online: http://www.businessweek.com/careers/content/jan1990/b3677014.htm.

tributed to natural resources or services. Even land rent is mainly a locational premium, not a payment for resources *in situ* or natural services—one more piece of evidence that the flow of low entropy from nature is treated as a free good. If we think of two social classes struggling to divide the pie, we have laborers getting 70% and capitalists, business owners, and landowners together getting about 30%. This division represents a kind of balance of power in the social struggle. Neither side wants to include nature as a participant in production, which would require that nature's services be paid according to their scarcity and productivity.

Even if one wanted to pay for nature's contribution, who would collect on nature's behalf? There is no social class analogous to labor or capital that has an interest in seeing to it that nature's services are properly accounted and paid for. Historically the landlord class may to some extent have played the role of defender of nature's services, but that class hardly exists anymore, and few lament its demise. The government is the biggest landholder in the U.S., and it has followed a policy of cheap resources in order to benefit and ease the tensions between labor and capital. The existing classes, labor and capital, see it in their mutual interest not to share with a third party. Since in reality there is no third party, all that would be necessary is to pay into a fund a scarcity rent for natural resources and services, and then redistribute the fund back to labor and capital, perhaps on the same 70-30 division. This would get the cost accounting and prices right, and improve the efficiency of allocation, without necessarily affecting the distribution. Alternatively, since many of the goods and services provided by natural capital are nonmarket goods, the scarcity rent could go toward supplying other nonmarket goods. The government could either do this directly or could subsidize the private production of such goods. The rent could also be redistributed progressively by financing the abolition of regressive taxes.

Consequences of Distribution for Community and Health

The existing distribution of wealth is not only a precondition for efficient allocation; it is also a fundamental dimension of justice in society. As such, it affects us more directly than we might at first think. Evidence indicates that inequality of income distribution (independently of absolute poverty) has a substantial effect on rates of morbidity and mortality.[5] The relatively poor have higher incidences of death and sickness than the relatively rich, regardless of the absolute level of income of the relatively poor. The main reason investigators suggest is the extra stress associated with being

[5]See. G. Wilkinson, *Mind the Gap*, New Haven, CT: Yale University Press, 2001.

relatively poor, being at the bottom of a dominance hierarchy. This extra stress is caused by less control over the circumstances of one's life, greater risks of job loss, a lower level of social standing and respect, more frequent experiences of disrespect and shame, with consequent anger and violence. Life at the bottom is more threatening, and the threat often comes from stressful relations with people higher up, including bosses, landlords, and government officials. Stress, of course, has well-known negative direct physiological effects on health.

In addition to these direct effects, inequality has indirect social effects on health. It is more difficult to form friendships across wider income gaps, as well as more difficult to form civic associations when wealth levels and economic interests are very disparate. Lack of friends and civic cohesion is also correlated with ill health. Treating people as atomistic, isolated individuals, unaffected by social relationships, literally makes them sick. As seen from our discussion of Max-Neef's human needs matrix (see Table 13.1), we are persons-in-community, related to each other internally—that is, our personal identity is largely constituted by our relation to others in the community. We are not independently defined entities held together only by external relations of the cash nexus. When these identity-constituting social connections become strained and corrupted by excessive inequality, we get sick more often and we die younger. We are also less happy.

Intertemporal Distribution of Wealth

Every bit as important as the distribution of wealth and income within a generation is the distribution of resources between generations. However, while people have pondered the distribution of resources within a generation for millennia, the concern for distribution between generations is more recent. For the vast majority of human history, natural resources appeared limitless and technological advance was slow. People had approximately the same resource endowment as their great grandparents had enjoyed, and they expected their great-grandchildren to inherit the same endowment as well. As the pace of technological change accelerated with the Industrial Revolution, change became noticeable from one generation to the next, and people began to expect a better life for their children than they themselves had enjoyed. The "Protestant work ethic" asked people to work hard and invest for their children. At least up through the 1960s, the question most economists asked was: How much consumption should this generation sacrifice for the ever-growing well-being of the next?[6]

[6]E.g., J. Robinson, *Essays in the Theory of Economic Growth*, London: Macmillan, 1968; E. Phelps, Second Essay on the Golden Rule of Accumulation, *American Economic Review*: 793–814 (1965).

However, the onset of the atomic age made it apparent that technological advance had the capacity to bring harm as well as good. Growth in population and per-capita consumption raised the specter of resource depletion. Worsening pollution caused alarm, and ecologists began to worry that many systems were nearing irreversible, catastrophic thresholds. The relevant question was no longer: How much should we sacrifice to make the future even better off? Now it was: How much should we sacrifice to keep the future from being worse off than the present? Curiously, at least in the United States, a culture change was occurring at about the same time. The work ethic was no longer "work hard, live frugally and invest in the future" but rather "work hard and consume as much as possible now." As a result, savings rates in the U.S. are currently at historic lows and rapidly approaching zero. For much of the population, savings rates are negative.

Should people strive to make the future better off than the present? Do we have at least an obligation to make sure it is not worse off than the present? There are no easy answers to the "appropriate" distribution of wealth between generations. Even a brief survey of philosophies is beyond the scope of this text. We will, however, quickly examine two alternative approaches: the ecological economics approach, based on ethical judgments concerning obligations to future generations (intergenerational justice), and the more mainstream approach in economics that argues for an "objective" decision-making rule (intergenerational allocation).

The Normative Approach of Ecological Economics

Ecological economists generally take the position that intergenerational resource distribution is an ethical issue. The generation into which someone is born is entirely based on chance. There is therefore no moral justification for claiming that one generation has any more right to natural resources, the building blocks of the economy, than any other. At the very least, future generations have an inalienable right to sufficient resources to provide a satisfactory quality of life. The current generation thus has a corresponding duty to preserve an adequate amount of resources. What is adequate depends on both technological and ecological change, both of which are characterized by pure uncertainty (ignorance). How we choose to deal with uncertainty is also an ethical decision.

What does this mean in practical terms?

Renewable and nonrenewable resources are fundamentally different and must be treated separately. An equal distribution of finite nonrenewable resources among a virtually infinite number of future generations would imply no resource use by any single generation. But there is no point in leaving resources in the ground forever, never to do anyone any good, so an upper limit to exhaustible resources for any one generation might be determined by the waste absorption capacity of the environment. As long as the use of the resource generates waste no faster than the ecosystem can absorb it, the use of exhaustible resources by one generation will not reduce renewable natural capital. Keeping fossil fuel use within such limits would automatically limit our ability to extract other mineral resources.

Even with a limited ability to extract nonrenewable but recyclable resources, each generation would have a further obligation to either efficiently recycle such resources, or at least minimize the generation and dispersion of "garbo-junk" as much as possible to make such resources as intergenerationally nonrival as possible. If existing technologies makes our well-being dependent on nonrenewable resources—as is currently the case—then we are simply obliged to develop substitutes for these resources. One option would be to capture marginal user costs, the unearned income from nonrenewable resources, and invest them toward developing such substitutes.[7]

Renewable resources as fund-services provide essential life-support functions, and these functions clearly must be maintained. Renewable resources as stock-flows must also be harvested at sustainable levels. No one created renewable resources, and therefore no single generation has the right to reduce the amount of the resource a future generation can sustainably consume, suggesting resource stocks must be at least as large as that which provides the maximum sustainable yield. As we saw in Chapter 12, sustainable management of renewable resources in a manner that "optimizes" both stock-flow and fund-service benefits will in general maintain these resources far from any catastrophic ecological thresholds. It is worth bearing in mind that as nonrenewable resources are finite, the exhaustion of these resources is a finite loss to future generations. Renewable resources, as both stock-flows and fund-services, produce a finite flow over an immeasurable number of future generations, and their irreversible loss therefore imposes a perpetual cost to the future.

The "Positive" Approach of Neoclassical Economics

Conventional economists, in contrast, favor an objective decision rule to determine the intergenerational allocation of resources. The problem thus becomes simply a technical one of comparing future benefits and costs with those that occur in the present. Fortunately, the market can tell us the value of things in the future relative to things today, and therefore the market can solve the problem of intergenerational allocation.

[7]For practical guidelines on investing scarcity rents, see S. El Serafy, "The Proper Calculation of Income from Depletable Natural Resources." In Y. J. Ahmad, S. El Serafy, and E. Lutz, eds., Washington, DC: *Environmental Accounting for Sustainable Development*, World Bank, 1989.

Intertemporal Discounting. How does the market reveal future values? Where adequate financial markets exist, people can borrow money today at interest, which requires them to pay back more money in the future. The fact that people engage in this activity reveals that people prefer things now rather than in the future, and economics must respect people's preferences.

There are three basic reasons why people might prefer things now to things in the future. First, people may simply be impatient. Anyone who goes into interest-bearing debt to purchase something is willing to sacrifice a greater quantity in the future for a smaller quantity now. Some of this impatience may come from uncertainty—no one knows for sure if he or she will be alive tomorrow, so why not eat, drink, and be merry today? This rationale for discounting is known as the **pure time rate of preference (PTRP)**.

Second, for things that reproduce, it makes sense that a given quantity in the future would be worth less than a given quantity now. For example, a handful of seed corn now can become a bushel of marketable corn a few months from now, so if growing corn was risk-free and required no resources or effort, then a handful now would be at least as valuable as a bushel in a few months. Of course, growing corn is risky and requires land, labor, and resources. However, market goods (in this case, seed corn) can be sold for money. Investing the money earned from the sale of the seed corn in an insured bank is not very risky, and for the individual investor basically requires no further resources or labor. As we explained in Chapter 10, this rationale for discounting is known as **opportunity cost**, the lost opportunity to invest. If money is a substitute for any other resource, then we should give more weight to any resource today over the same resource tomorrow.

Third, the economy has grown fairly steadily for hundreds of years. People therefore expect that they will be richer in the future than they are today. Just as an extra $1000 provides less utility to Bill Gates than to a pauper, the law of diminishing marginal utility means that money in the future will be worth less than the same amount of money today. This is sometimes referred to as the "richer future" argument.

In general, this process of valuing the future less than the present is known as **intertemporal discounting**, introduced in Chapter 10. Business people explicitly discount the future when making investment decisions, and mainstream economists argue that people automatically apply this concept to all of their purchase decisions. As a result, they conclude, the market efficiently allocates goods between the present and future.

What's more, if intertemporal discounting leads to allocative efficiency in the market, then it should also be applied to nonmarket investments. For example, one of the biggest nonmarket decisions we face today is how to deal with global climate change. Virtually all economic analyses of climate change place a lower weight on future costs and benefits than on present ones. These analyses look at different policy scenarios and for each sum up the present costs and benefits with discounted future costs and benefits to arrive at a **net present value (NPV)**. NPV basically tells us what present and future costs and benefits are worth to us *today* (not to the future *tomorrow*), which implies that future generations have no particular right to any resources, and we have no obligation to preserve any. Under this type of benefit-cost analysis, the higher the NPV relative to required investments, the better the project.

Such analyses can carry a great deal of weight as society decides how to address some of the most pressing problems we now confront. The central importance of the discount rate in determining the outcome of such analyses means the topic deserves our attention.

Box 15-2 Intertemporal Discounting and Global Climate Change

Policy makers seeking an objective decision-making tool for resolving the problems of global climate change have turned to economists. Economists typically respond to the problem by creating complex models of future costs and benefits of climate change, and compare these to the costs of mitigation measures in a cost-benefit analysis designed to calculate net present value. Not surprisingly, analyses using a fairly high discount rate find that future damages from global warming do not justify efforts today to reduce greenhouse gases. The 6% discount rate used in one study would have us believe that we should not invest $300 million today to prevent $30 trillion (a rough estimate of today's global GNP) in damages in 200 years.[a] A similar study using a 2% discount rate in contrast finds that we should make substantial investments now to reduce the impacts of global warming in the future. Similarly, what we decide to do with an old-growth forest that supplies a small but steady flow of benefits forever if left intact or a large, one-time return if it is clear-cut will depend on the discount rate we choose.

A frequently asked question is: Does a higher or lower discount rate favor the environment? For a given fishery or mine, as we have seen, higher discount rates increase the intensity and rate of exploitation and are thus bad for the environment. But a higher interest rate (discount rate) slows down aggregate growth in GNP and throughput, thus easing pressure on the environment. In terms of evaluating a given project, a high discount rate favors projects whose costs are mainly in the future and whose benefits are in the present, and penalizes those who costs are in the present with benefits in the future. Most issues in economics are not simple, and that certainly holds for discounting.

In many such models, the choice of a discount rate may be the single most important factor, yet respected economists addressing the same problem use dramatically different rates and arrive at dramatically different results. Are such models actually objective decision-making tools?

[a]30 trillion is a number that's hard to wrap your mind around. Putting it into perspective: 30 trillion seconds is slightly less than a million years (951,294 years, to be precise)!

Discounting Reconsidered

We have already explained why intertemporal discounting can make sense for the individual and for market goods. We must now examine whether it also makes sense for society and for nonmarket goods.

We saw why individuals might have a pure time rate of preference: People are impatient; they don't live forever; possessions can be lost, destroyed, or stolen, and opportunities disappear. A reasonable individual may discount the future for any one of these reasons—why should I pay money now to reduce damages from global warming that will only occur after I am dead?—but the same logic does not apply to society. Relative to the individuals of which they are composed, societies are immortal, and uncertainties are averaged out. For this reason, there is, in fact, fairly wide consensus within the economics profession that social discount rates should indeed be lower than individual discount rates. The **social discount rate** is a rate of conversion of future value to present value that reflects society's collective ethical judgment, as opposed to an individualistic judgment, such as the market rate of interest.

When it comes to the opportunity cost of capital, however, the consensus changes. Financial capital does function as a productive asset, and if we have it now instead of in the future, we have the opportunity to invest it in productive activities that will increase the quantity of market goods in the future. There are a number of important issues we must bear in mind, however.

First, the real value of money can only grow if the production of goods and services that money can acquire also grows, and we know that the production of goods and services cannot grow forever on a finite planet. While there may always be some areas that are growing, justifying a discount rate for the individual, the economy as a whole cannot grow indefinitely, in which case a social discount rate into the indefinite future may be inappropriate.

Second, we must recognize that many investments are "profitable" because we ignore many of the costs of production. We know that all human productive activities use up natural resources and return waste to the environment, and these costs of production are often ignored. Many of these costs, such as contributions to global warming, have greater impacts on future generations. Thus, ignoring costs to future generations allows us to earn higher returns on investments. We then use these higher returns to justify the fact that we ignore costs imposed on future generations! Even in the short run, then, it seems that market-determined interest rates are not suitable discount rates.

Related to the opportunity cost of capital is the argument that the future will be richer than the present because of investments we make now. Of course, if the economy does not continue to grow, the future will not be richer, and if we deplete our natural resource stock, there is every chance the future will be poorer. In fact, measures such as the ISEW suggest that society is already growing poorer, not richer, if we take into account external costs. Also, if we believe that natural capital must be treated separately from manmade capital (because they are complements rather than substitutes and natural capital has become the limiting factor), then the decline in natural capital, coupled with the law of diminishing marginal utility, suggests we should apply a negative discount rate to natural capital. At the very least, we might consider applying a positive discount rate only to those goods and services that are actually highly fungible with money—that is, that can be converted into money and back again with little effort. Basically, this would mean that we should only discount market goods and services.

Third, there are only finite opportunities for productive investment in an economy; investments, like other things, show diminishing marginal returns. For example, someone borrows money to explore for oil, and someone else borrows money to build a car factory. The next person to borrow money to explore for oil will have fewer places to explore, and therefore will expect lower returns. The next person to borrow to build a car factory will face a more saturated market, and therefore expect to sell fewer cars. As more and more people borrow to invest and opportunities are used up, the returns on investments can be expected to decrease, ultimately falling to zero. More likely, if interest rates are determined in the market by the supply and demand of money for borrowing and returns on investments, a balance will be reached in which investors cannot afford to pay high enough interest rates for consumers to be interested in deferring consumption. Theoretically then, in a perfect market situation, the opportunity cost of capital at the margin will just equal the PTRP of the existing generation. (Obviously, future generations cannot take part in financial markets any more than they can in any other market.) However, a high PTRP means consumption will be high and investment and growth low,[8] and a low PTRP implies the op-

[8]Returns on investments will be high, but the total amount invested will be low, and hence growth will be low.

posite. Thus, if we allow the market interest rate to determine the discount rate, there would theoretically be an inverse correlation between discount rate and economic growth, the exact opposite of what would justify the "richer future" rationale for discounting.

Box 15-3 Discounting, Psychology, and Economics

Economists argue that economics is the science of human preferences, so human preferences must be respected. If people value the present more than the future, we must respect that. The question is: Do people *exponentially* discount the future? While it is true you may prefer to have something now rather than the same thing in 5 years, how do you value something that happens 100 years in the future compared to 105 years? If you heard that global warming was going to result in the deaths of 50 million Bangladeshis in 125 years, would that make you feel only half as bad as finding out it would actually kill those 50 million Bangladeshis in 100 years? If you are like most people, you would feel equally bad in either case, yet influential economic models of the impacts of global warming really do assume we would care only half as much about those deaths if they occurred 25 years later.

Empirical studies show that people do discount the future, but do not do so exponentially. We might give more weight to what happens now than to what happens in the near future, but we are nearly indifferent between the same outcome occurring at different times in the more distant future. One approach to modeling this type of behavior mathematically is known as **hyperbolic discounting**. While this precise formulation of intertemporal discounting may not be perfect, considerable evidence suggests it is far more representative of human preferences than exponential discounting. While the approach was first introduced over 30 years ago and has gained increasing attention in the last few years, it is still fairly rare to see it in use.

An increasing number of studies in the area of Psychology and Economics are finding that the traditional economic assumptions of human behavior are often seriously flawed. If economics is serious about becoming the science of human preferences, it would do well to pay more attention to how humans really behave.[a]

---

[a] *For a good introduction to the field of Psychology and Economics, see M. Rabin, Psychology and Economics,* Journal of Economic Literature *36(1): 11–46 (March 1998).*

Finally, many economists argue that technology is the driving force for economic growth. Not only does technology ensure we won't run out of resources and the economy won't stop growing, but it offers yet another reason to discount the value of resources in the future. Technology is likely to develop substitutes for natural resources. When these substitutes become available, the resources they replace will lose value. Therefore, they will be worth less in the future than they are today. After all, hasn't oil largely replaced coal, and haven't fiber optics replaced copper in many uses? However, technology ultimately complements resources and can never completely replace them. Some 150 years ago, oil had little value. Today, it is an integral part of an overwhelming number of industrial processes and products. As we saw in Chapter 5, we are actually developing new uses for oil and other raw materials faster than we are developing replacements, again suggesting that the value of raw materials will increase in the future, not shrink.

What in the end can we say about discount rates? They do make sense for individuals in the short run. For some small-scale, short-term social projects, they may also make sense. However, justifications for discounting the future on a large scale and over long-time horizons are questionable at best. **Intertemporal allocation** is the apportionment of resources across different stages in the lifetimes of basically the same set of people (the same generation). Discounting can make sense for someone efficiently allocating resources intertemporally. But as we lengthen the time period we are more and more talking about different people (different generations) and less and less about the same people at different stages in their lives. **Intertemporal distribution** is the apportioning of resources across different generations, across different people. Distribution is fundamentally different from allocation, and, consequently, justice replaces efficiency as the relevant criterion for policy when time periods become intergenerational.

BIG IDEAS to remember

- Pareto optimality
- Role of scale and distribution in defining Pareto optimal allocation
- Income distribution vs. wealth distribution
- Functional vs. personal distribution
- Social limits to range of inequality
- Lorenz curve
- Gini coefficient
- Inequality and health
- Intertemporal distribution vs. intertemporal allocation
- Discounting and net present value
- Pure time rate of preference (PTRP)
- Individual vs. social discount rates